
Inference via Interpolation: Contrastive Representations Provably Enable Planning and Inference

Benjamin Eysenbach*
Princeton University
eysenbach@princeton.edu

Vivek Myers*
UC Berkeley
vmyers@berkeley.edu

Ruslan Salakhutdinov
Carnegie Mellon University
rsalakhu@cs.cmu.edu

Sergey Levine
UC Berkeley
svlevine@eecs.berkeley.edu

Abstract

Given time series data, how can we answer questions like “what will happen in the future?” and “how did we get here?” These sorts of probabilistic inference questions are challenging when observations are high-dimensional. In this paper, we show how these questions can have compact, closed form solutions in terms of learned representations. The key idea is to apply a variant of contrastive learning to time series data. Prior work already shows that the representations learned by contrastive learning encode a probability ratio. By extending prior work to show that the marginal distribution over representations is Gaussian, we can then prove that joint distribution of representations is also Gaussian. Taken together, these results show that representations learned via temporal contrastive learning follow a Gauss-Markov chain, a graphical model where inference (e.g., prediction, planning) over representations corresponds to inverting a low-dimensional matrix. In one special case, inferring intermediate representations will be equivalent to interpolating between the learned representations. We validate our theory using numerical simulations on tasks up to 46-dimensions.¹

1 Introduction

Probabilistic modeling of time-series data has applications ranging from robotic control [84] to material science [42], from cell biology [74] to astrophysics [55]. These applications are often concerned with two questions: *predicting* future states (e.g., what will this cell look like in an hour), and *inferring* trajectories between two given states. However, answering these questions often requires reasoning over high-dimensional data, which can be challenging as most tools in the standard probabilistic toolkit require generation. Might it be possible to use discriminative methods (e.g., contrastive learning) to perform such inferences?

Many prior works aim to learn representations that are easy to predict while retaining salient bits of information. For time-series data, we want the representation to remain a sufficient statistic for distributions related to time – for example, they should retain bits required to predict future states (or representations thereof). While generative methods [26, 56, 101, 103] have this property, they tend to be computationally expensive (see, e.g., [72]) and can be challenging to scale to high-dimensional observations.

¹Code: https://github.com/vivekmyers/contrastive_planning

* Equal contribution.




Figure 1: We apply temporal contrastive learning to observation pairs to obtain representations $(\psi(x_0), \psi(x_{t+k}))$ such that $A\psi(x_0)$ is close to $\psi(x_{t+k})$. While inferring waypoints in the high-dimensional observation space is challenging, we show that the distribution over intermediate latent representations has a closed form solution corresponding to linear interpolation between the initial and final representations.

In this paper, we will study how contrastive methods (which are discriminative, rather than generative) can perform inference over time series. Ideally, we want representations of observations x to be a sufficient statistic for temporal relationships (e.g., does x' occur after x ?) but need not retain other information about x (e.g. the location of static objects). This intuition motivates us to study how contrastive representation learning methods [18, 65, 81, 87, 96] might be used to solve prediction and planning problems on time series data. While prior works in computer vision [16, 65] and natural language processing (NLP) [59] often study the geometry of learned representations, our results show how geometric operations such as interpolation are related to inference. Our analysis will focus on a regularized version of the symmetrized infoNCE objective [71], generating positive examples by sampling pairs of observations from the same time series data. We will study how representations learned in this way can facilitate two inference questions: prediction and planning.² As a stepping stone, we will build upon prior work [90] to show that regularized contrastive learning should produce representations whose marginal distribution is an isotropic Gaussian distribution.

The main contribution of this paper is to demonstrate how intermediate and future time steps in a time series can be inferred easily using contrastive representations. This inference problem captures a number of practical tasks: interpolation, in-filling, and even planning and control, where the intermediate steps represent states between a stand and goal. While ordinarily these problems require an iterative inference or optimization procedure, with contrastive representations this can be done simply by inverting a low-dimensional matrix. In one special case, inference will correspond to linear interpolation. Our first step is to prove that, under certain assumptions, the distribution over future representations has a Gaussian distribution, with a mean that is a linear function of the initial state representation (Lemma 1). This paves the way to our main result (Theorem 4.1): *given an initial and final state, we show that the posterior distribution over an intermediate state representations also follows a Gaussian distribution.* Said in other words, the representations follow a Gauss-Markov chain,³ wherein any joint or conditional distribution can be computed by inverting a low-dimensional matrix [57, 93] (See Fig. 1). In one special case, inference will correspond to linearly interpolating between the representations of an initial state and final state. Section 5 provides numerical experiments.

2 Related Work

Representations for time-series data. In applications ranging from robotics to vision to NLP, users often want to learn representations of observations from time series data such that the *spatial* arrangement of representations reflects the *temporal* arrangement of the observations [28, 59, 65, 70]. Ideally, these representations should retain information required to predict future observations and infer likely paths between pairs of observations. Many approaches use an autoencoder, learning rep-

²Following prior work [6, 10], we will use *planning* to refer to the problem of inferring intermediate states, not to refer to an optimal control problem.

³This probabilistic model is equivalent to a discretized Ornstein-Uhlenbeck process [88] and is also known as an AR(1) model [11, Eq. 3.1.16].

representations that retain the bits necessary to reconstruct the input observation, while also regularizing the representations to compressed or predictable [12, 20, 43, 45, 68, 104]. A prototypical method is the sequential VAE [101], which is computationally expensive to train because of the reconstruction loss, but is easy to use for inference. Our work shares the aims of prior methods that attempt to linearize the dynamics of nonlinear systems [7, 23, 63, 64, 80, 92], including videos [33, 40]. Our work aims to retain uncertainty estimates over predictions (like the sequential VAE) without requiring reconstruction. Avoiding reconstruction is appealing *practically* because it decreases the computational requirements and number of hyperparameters; and *theoretically* because it means that representations only need to retain bits about temporal relationships and not about the bits required to reconstruct the original observation.

Contrastive Learning. Contrastive learning methods circumvent reconstruction by learning representations that merely classify if two events were sampled from the same joint distribution [17, 35, 71]. When applied to representing states along trajectories, contrastive representations learn to classify whether two points lie on the same trajectory or not [29, 65, 70, 78, 97]. Empirically, prior work in computer vision and NLP has observed that contrastive learning acquires representations where interpolation between representations corresponds to changing the images in semantically meaningful ways [19, 50, 59, 67, 95, 98].

Our analysis will be structurally similar to prior theoretical analysis on explaining why word embeddings can solve analogies [3, 5, 48]. Our work will make a Gaussianity assumption similar to Arora et al. [5] and our Markov assumption is similar to the random walks analyzed in Arora et al. [5], Hashimoto et al. [36]. Our paper builds upon and extends these results to answer questions such as: “what is the distribution over future observations representations?” and “what is the distribution over state (representations) that would occur on the path between one observation and another?” While prior work is primarily aimed at explaining the good performance of contrastive word embeddings (see, e.g., [5]), we are primarily interested in showing how similar contrastive methods are an effective tool for inference over high-dimensional time series data. Our analysis will show how representations learned via temporal contrastive learning (i.e., without reconstruction) are sufficient statistics for inferring future outcomes and can be used for performing inference on a graphical model (a problem typically associated with generative methods).

Goal-oriented decision making. Much work on time series representations is done in service of learning goal-reaching behavior, an old problem [46, 61] that has received renewed attention in recent years [14, 15, 21, 37, 39, 52, 76, 99]. Some of the excitement in goal-conditioned RL is a reflection of the recent success of self-supervised methods in computer vision [73] and NLP [66]. Our analysis will study a variant of contrastive representation learning proposed in prior work for goal-conditioned RL [29, 78]. These methods are widespread, appearing as learning objectives for learning value functions [1, 28, 29, 49, 51, 53, 60, 86, 91, 102], as auxiliary objectives [4, 9, 13, 60, 77, 82, 83], in objectives for model-based RL [2, 32, 58, 80], and in exploration methods [25, 34]. Our analysis will highlight connections between these prior methods, the classic successor representation [8, 24], and probabilistic inference.

Planning. Planning lies at the core of many RL and control methods, allowing methods to infer the sequence of states and actions that would occur if the agent navigated from one state to a goal state. While common methods such as PRM [44] and RRT [47] focus on building random graphs, there is a strong community focusing on planning methods based on probabilistic inference [6, 85, 94]. The key challenge is scaling to high-dimensional settings. While semi-parametric methods make progress on this problem this limitation through semi-parametric planning [27, 30, 100], it remains unclear how to scale any of these methods to high-dimensional settings when states do not lie on a low-dimensional manifold. Our analysis will show how contrastive representations may lift this limitation, with experiments validating this theory on 39-dimensional and 46-dimensional tasks.

3 Preliminaries

Our aim is to learn representations of time series data such that the spatial arrangement of representations corresponds to the temporal arrangement of the underlying data: if one example occurs shortly after another, then they should be mapped to similar representations. This problem setting arises in many areas, including video understanding and reinforcement learning. To define this problem formally, we will define a Markov process with states x_t indexed by time t :⁴

⁴This can be extended to *controlled* Markov processes appending the previous action to the observations.

$p(x_{1:T} | x_0) = \prod_{t=0}^{T-1} p(x_{t+1} | x_t)$. The dynamics $p(x_{t+1} | x_t)$ tell us the immediate next state, and we can define the distribution over states t steps in the future by marginalizing over the intermediate states, $p_t(x_t | x_0) = \int p(x_{1:t} | x_0) dx_{1:t-1}$. A key quantity of interest will be the γ -discounted state occupancy measure, which corresponds to a time-averaged distribution over future states:

$$p_{t+}(x_{t+} = x) = (1 - \gamma) \sum_{t=0}^{\infty} \gamma^t p_t(x_t = x). \quad (1)$$

Contrastive learning. Our analysis will focus on applying contrastive learning to a particular data distribution. Contrastive learning [35, 65, 75] acquires representations using “positive” pairs (x, x^+) and “negative” pairs (x, x^-) . While contrastive learning typically learns just one representation, we will use two different representation for the two elements of the pair; that is, our analysis will use terms like $\phi(x)$, $\psi(x^+)$ and $\psi(x^-)$. We assume all representations lie in \mathbb{R}^k .

The aim of contrastive learning is to learn representations such that positive pairs have similar representations ($\phi(x) \approx \psi(x^+)$) while negative pairs have dissimilar representations ($\phi(x) \neq \psi(x^-)$). Let $p(x, x^+)$ be the joint distribution over positive pairs (i.e., $(x, x^+) \sim p(x, x^+)$). We will use the product of the marginal distributions to sample negative pairs ($(x, x^-) \sim p(x)p(x)$). Let B be the batch size, and note that the positive samples x_j^+ at index j in the batch serve as *negatives* for x_i for any $i \neq j$. Our analysis is based on the infoNCE objective without resubstitution [65, 81]:

$$\max_{\phi(\cdot), \psi(\cdot)} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\sum_{i=1}^B \log \frac{e^{-\frac{1}{2} \|\phi(x_i) - \psi(x_i^+)\|_2^2}}{\sum_{j \neq i} e^{-\frac{1}{2} \|\phi(x_i) - \psi(x_j^+)\|_2^2}} + \log \frac{e^{-\frac{1}{2} \|\phi(x_i) - \psi(x_i^+)\|_2^2}}{\sum_{j \neq i} e^{-\frac{1}{2} \|\phi(x_j) - \psi(x_i^+)\|_2^2}} \right] \quad (2)$$

We will use the symmetrized version of this objective [71], where the denominator is the sum across rows of a logits matrix and once where it is a sum across the.

While contrastive learning is typically applied to an example x and an augmentation $x^+ \sim p(x | x)$ of that same example (e.g., a random crop), we will follow prior work [65, 78] in using the time series *dynamics* to generate the positive pairs, so x^+ will be an observation that occurs temporally after x . While our experiments will sample positive examples from the discounted state occupancy measure ($x^+ \sim p_{t+}(x_{t+} | x)$) in line with prior work [29], our analysis will also apply to different distributions (e.g., always sampling a state k steps ahead).

While prior work typically constrains the representations to have a constant norm (i.e., to lie on the unit hypersphere) [65], we will instead constrain the *expected* norm of the representations is bounded, a difference that will be important for our analysis:

$$\frac{1}{k} \mathbb{E}_{p(x)} [\|\psi(x)\|_2^2] \leq c. \quad (3)$$

Because the norm scales with the dimension of the representation, we have scaled down the left side by the representation dimension, k . In practice, we will impose this constraint by adding a regularization term $\lambda \mathbb{E}_{p(x)} [\|\psi(x)\|_2^2]$ to the infoNCE objective (Eq. 2) and dynamically tuning the weight λ via dual gradient descent.

3.1 Key assumptions

This section outlines the two key assumptions behind our analysis, both of which have some theoretical justification. Our main assumption examines the distribution over representations:

Assumption 1. Regularized, temporal contrastive learning acquires representations whose marginal distribution representations $p(\psi) \triangleq \int p(x) \mathbb{1}(\psi(x) = \psi) dx$ is an isotropic Gaussian distribution:

$$p(\psi) = \mathcal{N}(\psi; \mu = 0, \sigma = c \cdot I). \quad (4)$$

In Appendix A.1 we extend prior work [90] provide some theoretical intuition for why this assumption should hold: namely, that the isotropic Gaussian is the distribution that maximizes entropy subject to an expected L2 norm constraint (Eq. 3) [22, 41, 79]. Our analysis also assumes that the learned representations converge to the theoretical minimizer of the infoNCE objective:

Assumption 2. Applying contrastive learning to the symmetrized infoNCE objective results in representations that encode a probability ratio:

$$e^{-\frac{1}{2} \|\phi(x_0) - \psi(x)\|_2^2} = \frac{p_{t+}(x_{t+} = x | x_0)}{p(x)C}. \quad (5)$$

This assumption holds under ideal conditions [54, 69] (see Appendix A.5),⁵ but we nonetheless call this an “assumption” because it may not hold in practice due to sampling and function approximation error. This assumption means the learned representations are sufficient statistics for predicting the probability (ratio) of future states: these representations must retain all the information pertinent to reasoning about *temporal* relationships, but need not retain information about the precise contents of the observations. As such, they may be much more compressed than representations learned via reconstruction.

Combined, these assumptions will allow us to express the distribution over sequences of representations as a Gauss-Markov chain. The denominator in Assumption 2, $p(x)$, may have a complex distribution, but Assumption 1 tells us that the distribution over *representations* has a simpler form. This will allow us to rearrange Assumption 2 to express the conditional distribution over representations as the product of two Gaussian likelihoods. Note that the left hand side of Assumption 2 already looks like a Gaussian likelihood.

4 Contrastive Representations Make Inference Easy

In this section, our main result will be to show how representations learned by (regularized) contrastive learning are distributed according to a Gauss-Markov chain, making it straightforward to perform inference (e.g., planning, prediction) over these representations. Our proof technique will combine (known) results about Gaussian distributions with (known) results about contrastive learning. We start by discussing an important choice of parametrization (Section 4.1) that facilitates prediction (Section 4.2) before presenting the main result in Section 4.3.

4.1 A Parametrization for Shared Encoders

This section describes the two encoders ($\psi(\cdot), \phi(\cdot)$) to compute representations of x and x^+ . While prior work in computer vision and NLP literature use the same encoder for both x and x^+ , this decision does not make sense for many time-series data as it would imply that our prediction for $p(x_t | x_0)$ is the same as our prediction for $p(x_0 | x_t)$. However, the difficulty of transiting from x_0 to x_t (e.g., climbing to the peak of a mountain) might be more difficult than the reverse (e.g., sledding down a mountain). Our proposed parametrization will handle this asymmetry.

We will treat the encoder $\psi(\cdot)$ as encoding the contents of the state. We will additionally learn a matrix A so that the function $\psi \mapsto A\psi$ corresponds to a (multi-step) prediction of the future representation. To map this onto contrastive learning, we will use $\phi(x) \triangleq A\psi(x)$ as the encoder for the initial state. One way of interpreting this encoder is as an additional linear projection applied on top of $\psi(\cdot)$, a design similar to those used in other areas of contrastive learning [17]. Once learned, we can use these encoders to answer questions about prediction (Section 4.2) and planning (Section 4.3).

4.2 Representations Encode a Predictive Model

Given an initial state x_0 , what states are likely to occur in the future? Answering this question directly in terms of high-dimensional states is challenging, but our learned representations provide a straightforward answer. Let $\psi_0 = \psi(x_0)$ and $\psi_{t+} = \psi(x_{t+})$ be random variables representing the representations of the initial state and a future state. Our aim is to estimate the distribution over these future representations, $p(\psi_{t+} | \psi_0)$. We will show that the learned representations encode this distribution.

Lemma 1. *Under the assumptions from Section 3, the distribution over representations of future states follows a Gaussian distribution with mean parameter given by the initial state representation:*

$$p(\psi_{t+} = \psi | \psi_0) = \mathcal{N}\left(\mu = \frac{c}{c+1}A\psi_0, \Sigma = \frac{c}{c+1}I\right). \quad (6)$$

⁵While the result of Ma and Collins [54] has $C(x)$ depending on x , the symmetrized version [71] removes the dependence on x .




Figure 2: A parametrization for temporal contrastive learning.




Figure 3: Predicting representations of future states.

The main takeaway here is that the distribution over future representations has a convenient, closed form solution. The representation norm constraint, c , determines the shrinkage factor $\frac{c}{c+1} \in [0, 1)$; highly regularized settings (small c) move the mean closer towards the origin and decrease the variance, as visualized in Fig. 3. Regardless of the constraint c , the predicted mean is a linear function $\psi \mapsto \frac{c}{c+1} A\psi$. The proof is in Appendix A.2. The proof technique is similar to that of the law of the unconscious statistician.

4.3 Planning over One Intermediate State

We now show how these representations can be used for a specific type of planning: given an initial state x_0 and a future state x_{t+} , infer the representation of an intermediate ‘‘waypoint’’ state x_w . The next section will extend this analysis to inferring the entire sequence of intermediate states. We assume $x_0 \rightarrow x_w \rightarrow x_{t+}$ form a Markov chain where $x_w \sim p(x_{t+} | x_0 = x_0)$ and $x_{t+} \sim p(x_{t+} | x_0 = x_w)$ are both drawn from the discounted state occupancy measure (Eq. 1). Let random variable $\psi_w = \psi(x_w)$ be the representation of this intermediate state. Our main result is that the posterior distribution over waypoint representations has a closed form solution in terms of the initial state representation and future state representation:

Theorem 4.1. *Under Assumptions 1 and 2, the posterior distribution over waypoint representations is a Gaussian whose mean and covariance are linear functions of the initial and final state representations:*

$$p(\psi_w | \psi_0, \psi_{t+}) = \mathcal{N}\left(\psi_w; \mu = \Sigma(A^T \psi_{t+} + A\psi_0), \Sigma^{-1} = \frac{c}{c+1} A^T A + \frac{c+1}{c} I\right).$$

The proof (Appendix A.3) uses the Markov property together with Lemma 1. The main takeaway from this lemma is that the posterior distribution takes the form of a simple probability distribution (a Gaussian) with parameters that are linear functions of the initial and final representations.

We give three examples to build intuition:

Example 1: $A = I$ and the c is very large (little regularization). Then, the covariance is $\Sigma^{-1} \approx 2I$ and the mean is the simple average of the initial and final representations $\mu \approx \frac{1}{2}(\psi_0 + \psi_{t+})$. In other words, the waypoint representation is the midpoint of the line $\psi_0 \rightarrow \psi_{t+}$.

Example 2: A is a rotation matrix and c is very large. Rotation matrices satisfy $A^T = A^{-1}$ so the covariance is again $\Sigma^{-1} \approx 2I$. As noted in Section 4.2, we can interpret $A\psi_0$ as a prediction of which representations will occur after ψ_0 . Similarly, $A^{-1}\psi_{t+} = A^T\psi_{t+}$ is a prediction of which representations will occur before ψ_{t+} . Theorem 4.1 tells us that the mean of the waypoint distribution is the simple average of these two predictions, $\mu \approx \frac{1}{2}(A^T\psi_{t+} + A\psi_0)$.

Example 3: A is a rotation matrix and $c = 0.01$ (very strong regularization). In this case $\Sigma^{-1} = \frac{0.01}{0.01+1} A^T A + \frac{0.01+1}{0.01} I \approx 100I$, so $\mu \approx \frac{1}{100}(\psi_0 + \psi_{t+}) \approx 0$. Thus, in the case of strong regularization, the posterior concentrates around the origin.

4.4 Planning over Many Intermediate States

This section extends the analysis to multiple intermediate states. Again, we will infer the posterior distribution of the representations of these intermediate states, $\psi_{w_1}, \psi_{w_2}, \dots$. We assume that these states form a Markov chain.

Theorem 4.2. *Given observations from a Markov chain $x_0 \rightarrow x_1 \dots x_{t+}$, the joint distribution over representations is a Gaussian distribution. Using $\psi_{1:n} = (\psi_{w_1}, \dots, \psi_{w_n})$ to denote the concatenated representations of each observation, we can write this distribution as*

$$p(\psi_{1:n}) \propto \exp\left(-\frac{1}{2}\psi_{1:n}^T \Sigma^{-1} \psi_{1:n} + \eta^T \psi_{1:n}\right),$$

where Σ^{-1} is a tridiagonal matrix

$$\Sigma^{-1} = \begin{pmatrix} \frac{c}{c+1} A^T A + \frac{c+1}{c} I & & & & \\ & -A & & & \\ & & \frac{c}{c+1} A^T A + \frac{c+1}{c} I & & \\ & & & -A & \\ & & & & \ddots \end{pmatrix} \quad \text{and} \quad \eta = \begin{pmatrix} A\psi_0 \\ 0 \\ \vdots \\ A^T\psi_{t+} \end{pmatrix}.$$

This distribution can be written in the canonical parametrization as $\Sigma = \Lambda^{-1}$ and $\mu = \Sigma\eta$. Recall that Gaussian distributions are closed under marginalization. Thus, once in this canonical parametrization,

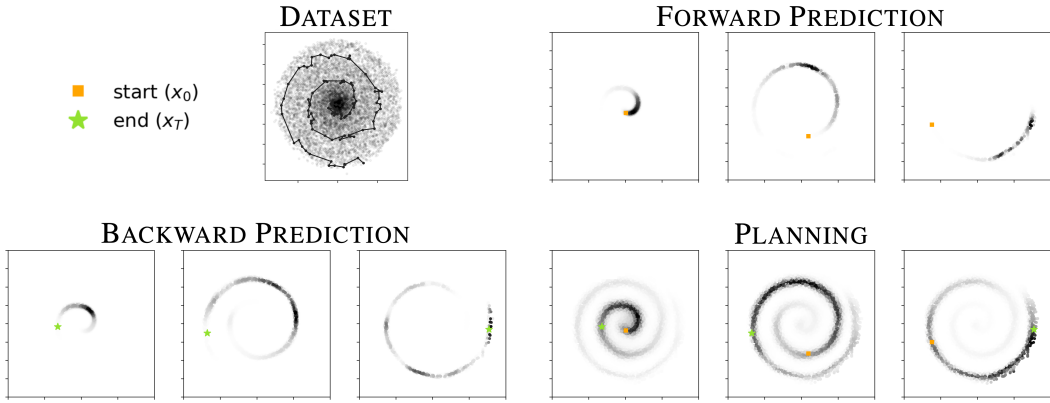


Figure 4: **Numerical simulation of our analysis.** (*Top Left*) Toy dataset of time-series data consisting of many outwardly-spiraling trajectories. We apply temporal contrastive learning to these data. (*Top Right*) For three initial observations (■), we use the learned representations to predict the distribution over future observations. Note that these distributions correctly capture the spiral structure. (*Bottom Left*) For three observations (★), we use the learned representations to predict the distribution over preceding observations. (*Bottom Right*) Given an initial and final observation, we plot the inferred posterior distribution over the waypoint (Section 4.3). The representations capture the shape of the distribution.

the marginal distributions can be obtained by reading off individual entries of these parameters:

$$p(\psi_i | \psi_0, \psi_{t+}) = \mathcal{N}\left(\psi_i; \mu_i = (\Sigma\eta)^{(i)}, \Sigma_i = (\Lambda^{-1})^{(i,i)}\right).$$

The key takeaway here is that this posterior distribution over waypoints is Gaussian, and it has a closed form expression in terms of the initial and final representations (as well as regularization parameter c and the learned matrix A).

In the general case of n intermediate states, the posterior distribution is

$$p(\psi_{w_1} \cdots \psi_{w_n} | \psi_0, \psi_{t+}) \propto e^{-\frac{1+\frac{1}{c}}{2} \sum_{i=1}^n \|\frac{c}{c+1} A\psi_{w_i} - \psi_{w_{i+1}}\|_2^2},$$

where $\psi_{w_0} = \psi_0$ and $\psi_{w_{n+1}} = \psi_{t+}$. This corresponds to a chain graphical model with edge potentials $f(\psi, \psi') = e^{-\frac{1+\frac{1}{c}}{2} \|\frac{c}{c+1} A\psi - \psi'\|_2^2}$.

Special case. To build intuition, consider the special case where A is a rotation matrix and c is very large, so $\frac{c}{c+1} A^T A + \frac{c+1}{c} \approx 2I$. In this case, Σ^{-1} is a (block) second difference matrix [38]:

$$\Sigma^{-1} = \begin{pmatrix} 2I & -I & & \\ -I & 2I & -I & \\ & & \ddots & \ddots \end{pmatrix}.$$

The inverse of this matrix has a closed form solution [62, Pg. 471], allowing us to obtain the mean of each waypoint in closed form:

$$\mu_i = (1 - \lambda(i))A\psi_0 + \lambda(i)A^T\psi_{t+}, \quad (7)$$

where $\lambda(i) = \frac{i}{n+1}$. Thus, each posterior mean is a convex combination of the (forward prediction from the) initial representation and the (backwards prediction from the) final representation. When A is the identity matrix, the posterior mean is simple linear interpolation between the initial and final representations!

5 Numerical Simulation

We include several didactic experiments to illustrate our results. All results and figures can be reproduced by running `make` in the source code: https://github.com/vivekmyers/contrastive_planning. The expected compute time is a few hours on a A6000 GPU. Figures in this section show error across different training and dataset split seeds.

5.1 Synthetic Dataset

To validate our analysis, we design a time series task with 2D points where inference over intermediate points (i.e., in-filling) requires nonlinear interpolation. Fig. 4 (*Top Left*) shows the dataset of time




Figure 5: Using inferred paths over our contrastive representations for control boosts success rates by $4.5\times$ on the most difficult goals ($18\% \rightarrow 84\%$). Alternative representation learning techniques fail to improve performance when used for planning.

series data, starting at the origin and spiraling outwards, with each trajectory using a randomly-chosen initial angle. We applied contrastive learning with the parametrization in Section 4.1 to these data and used the learned representations to solve prediction and planning problems (see Fig. 4 for details). Note that these predictions correctly handle the nonlinear structure of these data – states nearby the initial state in Euclidean space that are not temporally adjacent are assigned low likelihood.

5.2 Solving Mazes with Inferred Representations

Our next experiment studies whether the inferred representations are useful for solving a control task. We took a 2d maze environment and dataset from prior work (Fig. 5, Left) [31] and learned encoders from this dataset. To solve the maze, we take the observation of the starting state and goal state, compute the representations of these states, and use the analysis in Section 4.3 to infer the sequence of intermediate representations. We visualize the results using a nearest neighbor retrieval (Fig. 5, Left). Appendix Fig. 7 contains additional examples.

Finally, we studied whether these representations are useful for control. We implemented a simple proportional controller for this maze. As expected, this proportional controller can successfully navigate to close goals, but fails to reach distant goals (Fig. 5, Right). However, if we use the proportional controller to track a series of waypoints planned using our representations (i.e., the orange dots shown in Fig. 5 (Left)), the success rate increases by up to $4.5\times$. To test the importance of *nonlinear* representations, we compare with a “PCA” baseline that predicts waypoints by interpolating between the principal components of the initial state and goal state. The better performance of our method indicates the importance of doing the interpolation using representations that are *nonlinear* functions of the input observations. While prior methods learn representations to encode temporal distances, it is unclear whether these methods support inference via interpolation. To test this hypothesis, we use one of these methods (“VIP” [51]) as a baseline. While the VIP representations likely encode similar bits as our representations, the better performance of the contrastive representations indicates that the VIP representations do not expose those bits in a way that makes planning easy.

5.3 Higher dimensional tasks

In this section we provide preliminary experiments showing the planning approach in Section 4 scales to higher dimensional tasks. We used two datasets from prior work [31]: `door-human-v0` (39-dimensional observations) and `hammer-human-v0` (46-dimensional observations). After learning encoders on these tasks, we evaluated the inference capabilities of the learned representations. Given the first and last observation from a trajectory in a validation set, we use linear interpolation (see Eq. (7)) to infer the representation of five intermediate waypoint representations.

We evaluate performance in two ways. **Quantitatively**, we measure the mean squared error between each of the true waypoint observations and those inferred by our method. Since our method infers representations, rather than observations, we use a nearest-neighbor retrieval on a validation set so that

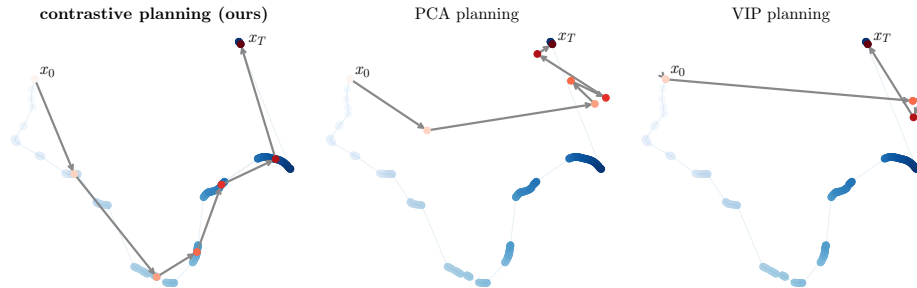




Figure 6: Planning for 39-dimensional robotic door opening. (Top Left) We use a dataset of trajectories demonstrating door opening from prior work [31] to learn representations. (Top Right) We use our method and three baselines to infer one intermediate waypoint between the first and last observation in a trajectory from a held-out validation set. Errors are measured using the mean squared error with the true waypoint observation; predicted representations are converted to observations using nearest neighbors on a validation set. (Bottom) We visualize a TSNE [89] of the states along the sampled trajectory as blue circles, with the transparency indicating the index along the trajectory. The inferred plan is shown as red circles connected by arrows. Our method generates better plans than alternative representation learning methods (PCA, VIP).

we can measure errors in the space of observations. **Qualitatively**, we visualize the high-dimensional observations from the validation trajectory using a 2-dimensional TSNE [89] embedding, overlying the infer waypoints from our method; as before, we convert the representations inferred by our method to observations using nearest neighbors.

We compare with three alternative methods in Fig. 6. To test the importance of representation learning, we first naïvely interpolate between the initial and final observations (“no planning”). The poor performance of this baseline indicates that the input time series are highly nonlinear. Similarly, interpolating the principle components of the initial and final observations (“PCA”) performs poorly, again highlighting that the input time series is highly nonlinear and that our representations are doing more than denoising (i.e., discarding directions of small variation). The third baseline, “VIP” [51], learns representations to encode temporal distances using approximate dynamic programming. Like our method, VIP avoids reconstruction and learns nonlinear representations of the observations. However, the results in Fig. 6 highlight that VIP’s representations do not allow users to plan by interpolation. The error bars shown in Fig. 6 (Top Right) show the standard deviation over 500 trajectories sampled from the validation set. For reproducibility, we repeated this entire experiment on another task, the 46-dimensional hammer-human-v0 from D4RL. The results, shown in Appendix Fig. 8, support the conclusions above. Taken together, these results show that our procedure for interpolating contrastive representations continues to be effective on tasks where observations have dozens of dimensions.

6 Discussion

Representation learning is at the core of many high-dimensional time-series modeling questions, yet how those representations are learned is often disconnected with the inferential task. The main contribution of this paper is to show how *discriminative* techniques can be used to acquire compact representations that make it easy to answer inferential questions about time. The precise objective and parametrization we studied is not much different from that used in practice, suggesting that either our theoretical results might be adapted to the existing methods, or that practitioners might adopt these details so they can use the closed-form solutions to inference questions. Our work may also have implications for studying the structure of learned representations. While prior work often studies

the geometry of representations as a post-hoc check, our analysis provides tools for studying *when* interpolation properties are guaranteed to emerge, as well as *how* to learn representations with certain desired geometric properties.

Limitations. Our analysis hinges on the two assumptions mentioned in Section 3.1, and it remains open how errors in those approximations translate into errors in our analysis. One important open question is whether it is always possible to satisfy these assumptions using sufficiently-expressive representations.

Acknowledgements. We thank Seohong Park, Gautam Reddy, Chongyi Zheng, and anonymous reviewers for feedback and discussions that shaped this project. This work was supported by Princeton Research Computing resources at Princeton University. This work was partially supported by AFOSR FA9550-22-1-0273.

References

- [1] Alekh Agarwal, Nan Jiang, Sham M Kakade, and Wen Sun. Reinforcement learning: Theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep*, pages 10–4, 2019.
- [2] Cameron S. Allen. Learning markov state abstractions for deep reinforcement learning. In *Neural Information Processing Systems*, 2021.
- [3] Carl Allen and Timothy Hospedales. Analogies explained: Towards understanding word embeddings. In *International Conference on Machine Learning*, pages 223–231. PMLR, 2019.
- [4] Ankesh Anand, Evan Racah, Sherjil Ozair, Yoshua Bengio, Marc-Alexandre Côté, and R Devon Hjelm. Unsupervised state representation learning in atari. *Advances in neural information processing systems*, 32, 2019.
- [5] Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. A latent variable model approach to pmi-based word embeddings. *Transactions of the Association for Computational Linguistics*, 4:385–399, 2016.
- [6] Hagai Attias. Planning by Probabilistic Inference. In *International Workshop on Artificial Intelligence and Statistics*, pages 9–16. PMLR, January 2003.
- [7] Ershad Banijamali, Rui Shu, Hung Bui, Ali Ghodsi, et al. Robust locally-linear controllable embedding. In *International Conference on Artificial Intelligence and Statistics*, pages 1751–1759. PMLR, 2018.
- [8] André Barreto, Will Dabney, Rémi Munos, Jonathan J Hunt, Tom Schaul, Hado P van Hasselt, and David Silver. Successor features for transfer in reinforcement learning. *Advances in neural information processing systems*, 30, 2017.
- [9] Homanga Bharadhwaj, Mohammad Babaeizadeh, Dumitru Erhan, and Sergey Levine. Information prioritization through empowerment in visual model-based rl. In *International Conference on Learning Representations*, 2021.
- [10] Matthew Botvinick and Marc Toussaint. Planning as inference. *Trends in cognitive sciences*, 16(10): 485–488, 2012.
- [11] George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- [12] Micah Carroll, Orr Paradise, Jessy Lin, Raluca Georgescu, Mingfei Sun, David Bignell, Stephanie Milani, Katja Hofmann, Matthew Hausknecht, Anca Dragan, et al. UniMASK: Unified Inference in Sequential Decision Problems, November 2022.
- [13] Pablo Samuel Castro, Tyler Kastner, P. Panangaden, and Mark Rowland. Mico: Improved representations via sampling-based state similarity for markov decision processes. In *Neural Information Processing Systems*, 2021.
- [14] Elliot Chane-Sane, Cordelia Schmid, and Ivan Laptev. Goal-conditioned reinforcement learning with imagined subgoals. In *International Conference on Machine Learning*, pages 1430–1440. PMLR, 2021.
- [15] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- [16] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [17] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15750–15758, 2021.

- [18] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020.
- [19] Ying-Cong Chen, Xiaogang Xu, Zhuotao Tian, and Jiaya Jia. Homomorphic Latent Space Interpolation for Unpaired Image-To-Image Translation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2403–2411, Long Beach, CA, USA, June 2019. IEEE. ISBN 978-1-72813-293-8. doi: 10.1109/CVPR.2019.00251.
- [20] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. A Recurrent Latent Variable Model for Sequential Data. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
- [21] Cédric Colas, Tristan Karch, Olivier Sigaud, and Pierre-Yves Oudeyer. Intrinsically motivated goal-conditioned reinforcement learning: a short survey. *preprint*, 2021.
- [22] Keith Conrad. Probability distributions and maximum likelihood, 2010. URL <https://api.semanticscholar.org/CorpusID:11225251>.
- [23] Brandon Cui, Yinlam Chow, and Mohammad Ghavamzadeh. Control-aware representations for model-based reinforcement learning. In *International Conference on Learning Representations*, 2020.
- [24] Peter Dayan. Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 5:613–624, 1993.
- [25] Yilun Du, Chuang Gan, and Phillip Isola. Curious representation learning for embodied intelligence. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10388–10397, 2021.
- [26] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Alex Lamb, Martin Arjovsky, Olivier Mastropietro, and Aaron Courville. Adversarially learned inference. In *International Conference on Learning Representations*, 2016.
- [27] Ben Eysenbach, Russ R Salakhutdinov, and Sergey Levine. Search on the Replay Buffer: Bridging Planning and Reinforcement Learning. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [28] Benjamin Eysenbach, Ruslan Salakhutdinov, and Sergey Levine. C-learning: Learning to achieve goals via recursive classification. In *International Conference on Learning Representations*, 2020.
- [29] Benjamin Eysenbach, Tianjun Zhang, Sergey Levine, and Russ R Salakhutdinov. Contrastive learning as goal-conditioned reinforcement learning. *Advances in Neural Information Processing Systems*, 35: 35603–35620, 2022.
- [30] Kuan Fang, Patrick Yin, Ashvin Nair, Homer Rich Walke, Gengchen Yan, and Sergey Levine. Generalization with lossy affordances: Leveraging broad offline data for learning visuomotor tasks. In *Conference on Robot Learning*, pages 106–117. PMLR, 2023.
- [31] Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4rl: Datasets for deep data-driven reinforcement learning. *arXiv preprint arXiv:2004.07219*, 2020.
- [32] Raj Ghugare, Homanga Bharadhwaj, Benjamin Eysenbach, Sergey Levine, and Russ Salakhutdinov. Simplifying model-based rl: Learning representations, latent-space models, and policies with one objective. In *The Eleventh International Conference on Learning Representations*, 2022.
- [33] Ross Goroshin, Michael F Mathieu, and Yann LeCun. Learning to linearize under uncertainty. *Advances in neural information processing systems*, 28, 2015.
- [34] Zhaohan Guo, Shantanu Thakoor, Miruna Pîslar, Bernardo Avila Pires, Florent Althé, Corentin Tallec, Alaa Saade, Daniele Calandriello, Jean-Bastien Grill, Yunhao Tang, et al. Byol-explore: Exploration by bootstrapped prediction. *Advances in neural information processing systems*, 35:31855–31870, 2022.
- [35] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, March 2010.
- [36] Tatsunori B Hashimoto, David Alvarez-Melis, and Tommi S Jaakkola. Word embeddings as metric recovery in semantic spaces. *Transactions of the Association for Computational Linguistics*, 4:273–286, 2016.
- [37] Joey Hejna, Jensen Gao, and Dorsa Sadigh. Distance weighted supervised learning for offline interaction data. *arXiv preprint arXiv:2304.13774*, 2023.
- [38] Nick Higham. What is the second difference matrix? <https://nhigham.com/2022/01/31/what-is-the-second-difference-matrix/>, 2022.
- [39] Michael Janner, Qiyang Li, and Sergey Levine. Offline reinforcement learning as one big sequence modeling problem. *Advances in neural information processing systems*, 34:1273–1286, 2021.
- [40] Dinesh Jayaraman and Kristen Grauman. Slow and steady feature analysis: higher order temporal coherence in video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3852–3861, 2016.

- [41] E. T. Jaynes. Information Theory and Statistical Mechanics. *Physical Review*, 106(4):620–630, May 1957. doi: 10.1103/PhysRev.106.620.
- [42] Hannes Jónsson, Greg Mills, and Karsten W Jacobsen. Nudged elastic band method for finding minimum energy paths of transitions. In *Classical and quantum dynamics in condensed phase simulations*, pages 385–404. World Scientific, 1998.
- [43] Siddharth Karamcheti, Suraj Nair, Annie S. Chen, Thomas Kollar, Chelsea Finn, Dorsa Sadigh, and Percy Liang. Language-Driven Representation Learning for Robotics, February 2023.
- [44] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE transactions on Robotics and Automation*, 12(4):566–580, 1996.
- [45] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.
- [46] John E Laird, Allen Newell, and Paul S Rosenbloom. Soar: An architecture for general intelligence. *Artificial intelligence*, 33(1):1–64, 1987.
- [47] Steven M LaValle, James J Kuffner, BR Donald, et al. Rapidly-exploring random trees: Progress and prospects. *Algorithmic and computational robotics: new directions*, 5:293–308, 2001.
- [48] Omer Levy and Yoav Goldberg. Linguistic regularities in sparse and explicit word representations. In *Proceedings of the eighteenth conference on computational natural language learning*, pages 171–180, 2014.
- [49] Bo Liu, Yihao Feng, Qiang Liu, and Peter Stone. Metric residual network for sample efficient goal-conditioned reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 8799–8806, 2023.
- [50] Xiaofeng Liu, Yang Zou, Lingsheng Kong, Zhihui Diao, Junliang Yan, Jun Wang, Site Li, Ping Jia, and Jane You. Data Augmentation via Latent Space Interpolation for Image Classification. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 728–733, August 2018. doi: 10.1109/ICPR.2018.8545506.
- [51] Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training. In *The Eleventh International Conference on Learning Representations*, 2022.
- [52] Yecheng Jason Ma, Jason Yan, Dinesh Jayaraman, and Osbert Bastani. How far i’ll go: Offline goal-conditioned reinforcement learning via f -advantage regression. *arXiv preprint arXiv:2206.03023*, 2022.
- [53] Yecheng Jason Ma, Vikash Kumar, Amy Zhang, Osbert Bastani, and Dinesh Jayaraman. LIV: Language-Image Representations and Rewards for Robotic Control. *International Conference on Machine Learning*, May 2023.
- [54] Zhuang Ma and Michael Collins. Noise contrastive estimation and negative sampling for conditional models: Consistency and statistical efficiency. In *Conference on Empirical Methods in Natural Language Processing*, 2018. URL <https://api.semanticscholar.org/CorpusID:52171904>.
- [55] Steven R Majewski, Ricardo P Schiavon, Peter M Frinchoy, Carlos Allende Prieto, Robert Barkhouser, Dmitry Bizyaev, Basil Blank, Sophia Brunner, Adam Burton, Ricardo Carrera, et al. The apache point observatory galactic evolution experiment (apogee). *The Astronomical Journal*, 154(3):94, 2017.
- [56] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- [57] Dmitry M Malioutov, Jason K Johnson, and Alan S Willsky. Walk-sums and belief propagation in gaussian graphical models. *The Journal of Machine Learning Research*, 7:2031–2064, 2006.
- [58] Bogdan Mazoure, Benjamin Eysenbach, Ofir Nachum, and Jonathan Tompson. Contrastive value learning: Implicit models for simple offline rl. In *Conference on Robot Learning*, pages 1257–1267. PMLR, 2023.
- [59] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [60] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A universal visual representation for robot manipulation. In *Conference on Robot Learning*, pages 892–909. PMLR, 2023.
- [61] Allen Newell, John C Shaw, and Herbert A Simon. Report on a general problem solving program. In *IFIP congress*, volume 256, page 64. Pittsburgh, PA, 1959.
- [62] Morris Newman and John Todd. The evaluation of matrix inversion programs. *Journal of the Society for Industrial and Applied Mathematics*, 6(4):466–476, 1958.
- [63] Tung Nguyen, Rui Shu, Tuan Pham, Hung Bui, and Stefano Ermon. Non-markovian predictive coding for planning in latent space, 2020.

- [64] Tung D Nguyen, Rui Shu, Tuan Pham, Hung Bui, and Stefano Ermon. Temporal predictive coding for model-based planning in latent space. In *International Conference on Machine Learning*, pages 8130–8139. PMLR, 2021.
- [65] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [66] OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023.
- [67] Alon Oring, Zohar Yakhini, and Yacov Hel-Or. Autoencoder image interpolation by shaping the latent space. In *International Conference on Machine Learning*, pages 8281–8290. PMLR, 2021.
- [68] Seongmin Park and Jihwa Lee. Finetuning Pretrained Transformers into Variational Autoencoders, November 2021.
- [69] Ben Poole, Sherjil Ozair, Aaron Van Den Oord, Alex Alemi, and George Tucker. On variational bounds of mutual information. In *International Conference on Machine Learning*, pages 5171–5180. PMLR, 2019.
- [70] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal Contrastive Video Representation Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6964–6974, 2021.
- [71] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [72] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019.
- [73] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [74] Wouter Saelens, Robrecht Cannoodt, Helena Todorov, and Yvan Saeys. A comparison of single-cell trajectory inference methods. *Nature biotechnology*, 37(5):547–554, 2019.
- [75] Nikunj Saunshi, Orestis Plevrakis, Sanjeev Arora, Mikhail Khodak, and Hrishikesh Khandeparkar. A theoretical analysis of contrastive unsupervised representation learning. In *International Conference on Machine Learning*, pages 5628–5637. PMLR, 2019.
- [76] Yannick Schroecker and Charles Isbell. Universal value density estimation for imitation learning and goal-conditioned reinforcement learning. *arXiv preprint arXiv:2002.06473*, 2020.
- [77] Max Schwarzer, Ankesh Anand, Rishab Goel, R Devon Hjelm, Aaron Courville, and Philip Bachman. Data-efficient reinforcement learning with self-predictive representations. In *International Conference on Learning Representations*, 2020.
- [78] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey Levine, and Google Brain. Time-contrastive networks: Self-supervised learning from video. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 1134–1141. IEEE, 2018.
- [79] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- [80] Rui Shu, Tung Nguyen, Yinlam Chow, Tuan Pham, Khoat Than, Mohammad Ghavamzadeh, Stefano Ermon, and Hung Bui. Predictive coding for locally-linear control. In *International Conference on Machine Learning*, pages 8862–8871. PMLR, 2020.
- [81] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems*, 29, 2016.
- [82] Adam Stooke, Kimin Lee, Pieter Abbeel, and Michael Laskin. Decoupling representation learning from reinforcement learning. In *International Conference on Machine Learning*, pages 9870–9879. PMLR, 2021.
- [83] Yunhao Tang, Zhaohan Daniel Guo, Pierre Harvey Richemond, Bernardo Avila Pires, Yash Chandak, Rémi Munos, Mark Rowland, Mohammad Gheshlaghi Azar, Charline Le Lan, Clare Lyle, et al. Understanding self-predictive learning for reinforcement learning. In *International Conference on Machine Learning*, pages 33632–33656. PMLR, 2023.
- [84] Evangelos Theodorou, Jonas Buchli, and Stefan Schaal. A generalized path integral control approach to reinforcement learning. *The Journal of Machine Learning Research*, 11:3137–3181, 2010.
- [85] Sep Thijssen and H. J. Kappen. Path integral control and state-dependent feedback. *Physical Review E*, 91(3):032104, March 2015. ISSN 1539-3755, 1550-2376. doi: 10.1103/PhysRevE.91.032104.
- [86] Stephen Tian, Suraj Nair, Frederik Ebert, Sudeep Dasari, Benjamin Eysenbach, Chelsea Finn, and Sergey Levine. Model-based visual planning with self-supervised functional distances. In *International*

- Conference on Learning Representations*, 2020.
- [87] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 776–794. Springer, 2020.
 - [88] George E Uhlenbeck and Leonard S Ornstein. On the theory of the brownian motion. *Physical review*, 36(5):823, 1930.
 - [89] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
 - [90] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International Conference on Machine Learning*, pages 9929–9939. PMLR, 2020.
 - [91] Tongzhou Wang, Antonio Torralba, Phillip Isola, and Amy Zhang. Optimal goal-reaching reinforcement learning via quasimetric learning. In *International Conference on Machine Learning*. PMLR, 2023.
 - [92] Manuel Watter, Jost Springenberg, Joschka Boedecker, and Martin Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. *Advances in neural information processing systems*, 28, 2015.
 - [93] Yair Weiss and William Freeman. Correctness of belief propagation in gaussian graphical models of arbitrary topology. *Advances in neural information processing systems*, 12, 1999.
 - [94] Grady Williams, Andrew Aldrich, and Evangelos Theodorou. Model Predictive Path Integral Control using Covariance Variable Importance Sampling, October 2015.
 - [95] Laurenz Wiskott and Terrence J. Sejnowski. Slow Feature Analysis: Unsupervised Learning of Invariances. *Neural Computation*, 14(4):715–770, April 2002. ISSN 0899-7667. doi: 10.1162/089976602317318938.
 - [96] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018.
 - [97] Mengda Xu, Zhenjia Xu, Cheng Chi, Manuela Veloso, and Shuran Song. Xskill: Cross embodiment skill discovery. In *Conference on Robot Learning*, pages 3536–3555. PMLR, 2023.
 - [98] Jia-Wei Yan, Ci-Siang Lin, Fu-En Yang, Yu-Jhe Li, and Yu-Chiang Frank Wang. Semantics-Guided Representation Learning with Applications to Visual Synthesis. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7181–7187, January 2021. doi: 10.1109/ICPR48806.2021.9413171.
 - [99] Rui Yang, Yiming Lu, Wenzhe Li, Hao Sun, Meng Fang, Yali Du, Xiu Li, Lei Han, and Chongjie Zhang. Rethinking goal-conditioned supervised learning and its connection to offline rl. In *International Conference on Learning Representations*, 2021.
 - [100] Tianjun Zhang, Benjamin Eysenbach, Ruslan Salakhutdinov, Sergey Levine, and Joseph E Gonzalez. C-planning: An automatic curriculum for learning goal-reaching tasks. In *International Conference on Learning Representations*, 2021.
 - [101] Shengjia Zhao, Jiaming Song, and Stefano Ermon. Towards deeper understanding of variational autoencoding models. *arXiv preprint arXiv:1702.08658*, 2017.
 - [102] Chongyi Zheng, Benjamin Eysenbach, Homer Rich Walke, Patrick Yin, Kuan Fang, Ruslan Salakhutdinov, and Sergey Levine. Stabilizing contrastive RL: Techniques for robotic goal reaching from offline data. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=Xkf2EBj4w3>.
 - [103] Yizhe Zhu, Martin Renqiang Min, Asim Kadav, and Hans Peter Graf. S3vae: Self-supervised sequential vae for representation disentanglement and data generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6538–6547, 2020.
 - [104] Yizhe Zhu, Martin Renqiang Min, Asim Kadav, and Hans Peter Graf. S3VAE: Self-Supervised Sequential VAE for Representation Disentanglement and Data Generation, May 2020.

A Proofs

A.1 Marginal Distribution over Representations is Gaussian

The infoNCE objective (Eq. (2)) can be decomposed into an alignment term and a uniformity term [90], where the uniformity term can be simplified as follows:

$$\begin{aligned}
& \mathbb{E}_{x \sim p(x)} \left[\log \mathbb{E}_{x^- \sim p(x)} \left[e^{-\frac{1}{2} \|A\psi(x^-) - \psi(x)\|_2^2} \right] \right] \\
&= \frac{1}{N} \sum_{i=1}^N \log \left(\frac{1}{N-1} \sum_{j=1 \dots N, j \neq i} e^{-\frac{1}{2} \|A\psi(x_i) - \psi(x_j)\|_2^2} \right) \\
&= \frac{1}{N} \sum_{i=1}^N \log \left(\frac{1}{N-1} \sum_{j=1 \dots N, j \neq i} \underbrace{\frac{1}{(2\pi)^{k/2}} e^{-\frac{1}{2} \|A\psi(x_i) - \psi(x_j)\|_2^2}}_{\mathcal{N}(\mu=\psi(x_j); \Sigma=I)} \right) + \frac{k}{2} \log(2\pi) \\
&= \frac{1}{N} \sum_{i=1}^N \log \hat{p}_{\text{GMM}}(\psi(x_i)) + \frac{k}{2} \log(2\pi) \\
&= -\hat{\mathcal{H}}[\psi(x)] + \frac{k}{2} \log(2\pi).
\end{aligned}$$

The derivation above extends that in Wang and Isola [90] by considering a Gaussian distribution rather than a von Mises Fisher distribution. We are implicitly making the assumption that the marginal distributions satisfy $p(x) = p(x^-)$. This difference corresponds to our choice of using a negative squared L2 distance in the infoNCE loss rather than an inner product, a difference that will be important later in our analysis. A second difference is that we do not use the resubstitution estimator (i.e., we exclude data point x_i from our estimate of \hat{p}_{GMM} when evaluating the likelihood of x_i), which we found hurt performance empirically. The takeaway from this identity is that maximizing the uniformity term corresponds to maximizing (an estimate of) the entropy of the representations.

We next prove that the maximum entropy distribution with an expected L2 norm constraint is a Gaussian distribution. Variants of this result are well known [22, 41, 79], but we include a full proof here for transparency.

Lemma 2. *The maximum entropy distribution satisfying the expected L2 norm constraint in Eq. (3) is a multivariate Gaussian distribution with mean $\mu = 0$ and covariance $\Sigma = c \cdot I$*

Proof. We start by defining the corresponding Lagrangian, with the second constraint saying that $p(x)$ must be a valid probability distribution.

$$\mathcal{L}(p) = \mathcal{H}_p[x] + \lambda_1 (\mathbb{E}_{p(x)} [\|x\|_2^2] - c \cdot k) + \lambda_2 \left(\int p(x) dx - 1 \right)$$

We next take the derivative w.r.t. $p(x)$:

$$\frac{\partial \mathcal{L}}{\partial p(x)} = -p(x)/p(x) - \log p(x) + \lambda_1 \|x\|_2^2 + \lambda_2$$

Setting this derivative equal to 0 and solving for $p(x)$, we get

$$p(x) = e^{-1 + \lambda_2 + \lambda_1 \|x\|_2^2}.$$

We next solve for λ_1 and λ_2 to satisfy the constraints in the Lagrangian. Note that $x \sim \mathcal{N}(\mu = 0, \Sigma = c \cdot I)$ has an expected norm $\mathbb{E}[\|x\|_2^2] = c \cdot k$, so we must have $\lambda_1 = -\frac{1}{2c}$. We determine λ_1 as the normalizing constant for a Gaussian, finally giving us:

$$p(x) = \frac{1}{(2c\pi)^{k/2}} e^{\frac{-1}{2c} \|x\|_2^2}$$

corresponding to an isotropic Gaussian distribution with mean $\mu = 0$ and covariance $\Sigma = c \cdot I$. \square

A.2 Proof of Lemma 1

Below we present the proof of Lemma 1

Proof. Our proof technique will be similar to that of the law of the unconscious statistician:

$$\begin{aligned}
p(\psi_{t+} | \psi_0) &\stackrel{(a)}{=} \frac{p(\psi_{t+}, \psi_0)}{p(\psi_0)} \propto \iint p(\psi_{t+}, x_{t+}, \psi_0, x_0) dx_{t+} dx_0 \\
&\stackrel{(b)}{=} \iint p(\psi_{t+} | x_{t+}) p(\psi_0 | x_0) p(x_{t+} | x_0) p(x_0) dx_{t+} dx_0 \\
&\stackrel{(c)}{=} \iint \mathbb{1}(\psi(x_{t+}) = \psi_{t+}) \mathbb{1}(\psi(x_0) = \psi_0) p(x_{t+}) e^{-\frac{1}{2} \|A\psi(x_0) - \psi(x_{t+})\|_2^2} p(x_0) dx_{t+} dx_0 \\
&\stackrel{(d)}{=} e^{-\frac{1}{2} \|A\psi_0 - \psi_{t+}\|_2^2} \iint \mathbb{1}(\psi(x_{t+}) = \psi_{t+}) \mathbb{1}(\psi(x_0) = \psi_0) p(x_{t+}) p(x_0) dx_{t+} dx_0 \\
&\stackrel{(e)}{=} e^{-\frac{1}{2} \|A\psi_0 - \psi_{t+}\|_2^2} \underbrace{\left(\int p(x_{t+}) \mathbb{1}(\psi(x_{t+})) dx_{t+} \right)}_{p(\psi_{t+})} \underbrace{\left(\int p(x_0) \mathbb{1}(\psi(x_0)) dx_0 \right)}_{p(\psi_0)} \\
&\stackrel{(f)}{\propto} e^{-\frac{1}{2} \|A\psi_0 - \psi_{t+}\|_2^2} e^{-\frac{1}{2c} \|\psi_{t+}\|_2^2} e^{-\frac{1}{2c} \|\psi_0\|_2^2} \\
&\stackrel{(g)}{\propto} e^{-\frac{1+\frac{1}{c}}{2} \left\| \frac{1}{1+\frac{1}{c}} A\psi_0 - \psi_{t+} \right\|_2^2} \\
&\propto \mathcal{N} \left(\psi_{t+}; \mu = \frac{c}{c+1} A\psi_0, \Sigma = \frac{c}{c+1} I \right).
\end{aligned}$$

In (a) we applied Bayes' Rule and removed the denominator, which is a constant w.r.t. ψ_{t+} . In (b) we factored the joint distribution, noting that ψ_{t+} and ψ_0 are deterministic functions of x_{t+} and x_0 respectively, so they are conditionally independent from the other random variables. In (c) we used Assumption 2 after solving for $p(x_{t+} | x_0) = p(x_{t+}) e^{-\frac{1}{2} \|A\psi(x_0) - \psi(x_{t+})\|_2^2}$. In (d) we noted that when the integrand is nonzero, it takes on a constant value of $e^{-\frac{1}{2} \|A\psi_0 - \psi_{t+}\|_2^2}$, so we can move that constant outside the integral. In (e) we used the definition of the marginal representation distribution (Eq. 6). In (f) we used Assumption 1 to write the marginal distributions $p(\psi_{t+})$ and $p(\psi_0)$ as Gaussian distributions. We removed the normalizing constants, which are independent of ψ_{t+} . In (g) we completed the square and then recognized the expression as the density of a multivariate Gaussian distribution. \square

A.3 Proof of Theorem 4.1: Waypoint Distribution

Proof.

$$\begin{aligned}
p(\psi_w | \psi_0, \psi_{t+}) &\stackrel{(a)}{=} \frac{p(\psi_{t+} | \psi_w) p(\psi_w | \psi_0)}{p(\psi_{t+} | \psi_0)} \\
&\stackrel{(b)}{\propto} e^{-\frac{1+\frac{1}{c}}{2} \left\| \frac{c}{c+1} A\psi_w - \psi_{t+} \right\|_2^2} e^{-\frac{1+\frac{1}{c}}{2} \left\| \frac{c}{c+1} A\psi_0 - \psi_w \right\|_2^2} \\
&\stackrel{(c)}{\propto} e^{-\frac{1}{2} (\psi_w - \mu)^T \Sigma^{-1} (\psi_w - \mu)} = \mathcal{N}(\psi_w; \mu, \Sigma)
\end{aligned}$$

where $\Sigma^{-1} = \frac{c}{c+1} A^T A + \frac{c+1}{c} I$ and $\mu = \Sigma(A^T \psi_{t+} + A\psi_0)$. \square

In line (a) we used the definition of the conditional distribution and then simplified the numerator using the Markov property. Line (b) uses the Lemma 1. Line (c) completes the square, the details of which are below:

$$\begin{aligned}
&\frac{1}{2} \cdot \frac{c+1}{c} \left(\left\| \frac{c}{c+1} A\psi_w - \psi_{t+} \right\|_2^2 + \left\| \frac{c}{c+1} A\psi_0 - \psi_w \right\|_2^2 \right) \\
&= \frac{1}{2} \cdot \frac{c+1}{c} \left(\psi_w^T \left(\frac{c}{c+1} A \right)^T \left(\frac{c}{c+1} A \right) \psi_w - 2\psi_{t+}^T \left(\frac{c}{c+1} A \right) \psi_w + \cancel{\psi_{t+}^T \psi_{t+}} \right. \\
&\quad \left. + \cancel{\psi_0^T \left(\frac{c}{c+1} A \right)^T \left(\frac{c}{c+1} A \right) \psi_0} - 2\psi_0^T \left(\frac{c}{c+1} A \right)^T \psi_w + \psi_w^T \psi_w \right)
\end{aligned}$$

$$\begin{aligned}
&\stackrel{\text{const.}}{=} \frac{1}{2} \frac{c+1}{c} \left(\psi_w^T \left(\left(\frac{c}{c+1} \right)^2 A^T A + I \right) \psi_w - 2 \frac{c}{c+1} (A^T \psi_{t+} + A \psi_0)^T \psi_w \right) \\
&= \frac{1}{2} \psi_w^T \underbrace{\left(\frac{c}{c+1} A^T A + \frac{c+1}{c} I \right)}_{\Sigma^{-1}} \psi_w - (A^T \psi_{t+} + A \psi_0)^T \psi_w \\
&\stackrel{\text{const.}}{=} (\psi_w - \mu)^T \Sigma^{-1} (\psi_w - \mu),
\end{aligned}$$

where $\Sigma^{-1} = \frac{c}{c+1} A^T A + \frac{c+1}{c} I$ and $\mu = \Sigma(A^T \psi_{t+} + A \psi_0)$. Above, we have used $\stackrel{\text{const.}}{=}$ to denote equality up to an additive constant that is independent of ψ_w .

A.4 Proof of Theorem 4.2: Planning over Many Intermediate States

Proof. We start by recalling that the waypoints form a Markov chain, so we can express their joint density as a product of conditional densities:

$$p(\psi_{1:n}) = p(\psi_0) p(\psi_1 | \psi_0) p(\psi_2 | \psi_1) \cdots p(\psi_n | \psi_{n-1}).$$

The aim of this lemma is to express the joint distribution over multiple waypoints, given an initial and final state representation:

$$\begin{aligned}
p(\psi_{1:n} | \psi_0, \psi_{t+}) &\stackrel{(a)}{=} \frac{p(\psi_{1:n} | \psi_0) p(\psi_{t+} | \psi_n)}{p(\psi_{t+} | \psi_0)} \\
&\stackrel{(b)}{\propto} p(\psi_1 | \psi_0) p(\psi_2 | \psi_1) \cdots p(\psi_{t+} | \psi_n) \\
&\stackrel{(c)}{\propto} \exp \left(-\frac{1}{2} \frac{c+1}{c} \|\frac{c}{c+1} A \psi_0 - \psi_1\|_2^2 - \frac{1}{2} \frac{c+1}{c} \|\frac{c}{c+1} A \psi_1 \right. \\
&\quad \left. - \psi_2\|_2^2 - \cdots - \frac{1}{2} \frac{c+1}{c} \|\frac{c}{c+1} A \psi_n - \psi_{t+}\|_2^2 \right) \\
&\stackrel{(d)}{=} \exp \left(-\frac{1}{2} \frac{c}{c+1} \psi_0^T A^T A \psi_0 + \psi_0^T A^T \psi_1 - \frac{1}{2} \frac{c+1}{c} \psi_1^T \psi_1 \right. \\
&\quad \left. - \frac{1}{2} \frac{c}{c+1} \psi_1^T A^T A \psi_1 + \psi_1^T A^T \psi_2 - \frac{1}{2} \frac{c+1}{c} \psi_2^T \psi_2 \right. \\
&\quad \left. - \frac{1}{2} \frac{c}{c+1} \psi_2^T A^T A \psi_2 + \psi_2^T A^T \psi_3 - \frac{1}{2} \frac{c+1}{c} \psi_3^T \psi_3 \right. \\
&\quad \quad \quad \vdots \\
&\quad \left. - \frac{1}{2} \frac{c}{c+1} \psi_n^T A^T A \psi_n + \psi_n^T A^T \psi_{t+} - \frac{1}{2} \frac{c+1}{c} \psi_{t+}^T \psi_{t+} \right) \\
&= \exp \left(-\frac{1}{2} \psi_{1:n}^T \Sigma^{-1} \psi_{1:n} + \eta^T \psi_{1:n} \right),
\end{aligned}$$

where Σ^{-1} is a tridiagonal matrix

$$\Sigma^{-1} = \begin{pmatrix} \frac{c}{c+1} A^T A + \frac{c+1}{c} I & -A^T & & & \\ -A & \frac{c}{c+1} A^T A + \frac{c+1}{c} I & -A^T & & \\ & -A & \frac{c}{c+1} A^T A + \frac{c+1}{c} I & & \\ & & & \ddots & \\ & & & & \ddots \end{pmatrix} \quad \text{and } \eta = \begin{pmatrix} A \psi_0 \\ 0 \\ \vdots \\ 0 \\ A^T \psi_{t+} \end{pmatrix}.$$

In (a) we applied Bayes' rule and removed the denominator because it is a constant with respect to $\psi_{1:n}$. In (b) we applied the Markov assumption. In (c) we used Lemma 1 to express the conditional probabilities as Gaussians, ignoring the proportionality constants (which are independent of ψ). In (d) we simplified the exponents, removing terms that do not depend on $\psi_{1:n}$. \square

A.5 Formalizing Assumption 2

Assumption 2 states

$$e^{-\frac{1}{2} \|\phi(x_0) - \psi(x)\|_2^2} = \frac{p(x | x_0)}{p(x^+) C} \quad (5)$$

when Eq. (2) is optimized.

We can justify this assumption by analyzing the general solution to the symmetrized version of the Oord et al. [65] infoNCE objective, which we do in Lemma 3. Applying this lemma to our representation learning objective (2) for sufficiently large batch size B then yields Eq. (5), with the function approximator $\|\phi(x) - \psi(x^+)\|^2 \approx f(x, x^+)$.

Lemma 3. *The solution to the optimization problem*

$$\max_{f(x, x^+)} \lim_{B \rightarrow \infty} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\frac{1}{B} \sum_{i=1}^B \log \frac{e^{f(x_i, x_i^+)}}{\sum_{j \neq i} e^{f(x_i, x_j^+)}} + \log \frac{e^{f(x_i, x_i^+)}}{\sum_{j \neq i} e^{f(x_j, x_i^+)}} \right] \quad (8)$$

satisfies

$$f(x, x^+) = \log \left(\frac{p(x^+ | x)}{p(x^+)C} \right) \quad (9)$$

for some C .

Proof of Lemma 3. We first break down the LHS and RHS of Eq. (2):

$$\begin{aligned} & \max_f \lim_{B \rightarrow \infty} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\underbrace{\frac{1}{B} \sum_{i=1}^B \log \frac{e^{f(x_i, x_i^+)}}{\sum_{j \neq i} e^{f(x_i, x_j^+)}}}_{\mathcal{J}_1} + \underbrace{\log \frac{e^{f(x_i, x_i^+)}}{\sum_{j \neq i} e^{f(x_j, x_i^+)}}}_{\mathcal{J}_2} \right] \\ \mathcal{J}_1(f) &= \lim_{B \rightarrow \infty} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\frac{1}{B} \sum_{i=1}^B \log \frac{e^{f(x_i, x_i^+)}}{\sum_{j \neq i} e^{f(x_i, x_j^+)}} \right] \\ \mathcal{J}_2(f) &= \lim_{B \rightarrow \infty} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\frac{1}{B} \sum_{i=1}^B \log \frac{e^{f(x_i, x_i^+)}}{\sum_{j \neq i} e^{f(x_j, x_i^+)}} \right] \end{aligned}$$

We now use the following result from Ma and Collins [54]:

Lemma 4. *The optimal solutions f_1 and f_2 for \mathcal{J}_1 and \mathcal{J}_2 satisfy*

$$f_1(x, x^+) = \log p(x | x^+) - \log c_1(x) \quad (10)$$

$$f_2(x, x^+) = \log p(x^+ | x) - \log c_2(x^+) \quad (11)$$

for arbitrary $c_1(x), c_2(x^+)$.

For any C , when $c_1(x) = Cp(x)$ and $c_2(x^+) = Cp(x^+)$,

$$f_1(x, x^+) = \log \left(\frac{p(x | x^+)}{p(x)C} \right) = \log \left(\frac{p(x^+ | x)}{p(x^+)C} \right) = f_2(x, x^+). \quad (12)$$

It follows that Eq. (12) maximizes both \mathcal{J}_1 and \mathcal{J}_2 , and is precisely the optimal solution Eq. (9) for Eq. (8). \square

What does C represent? From Eq. (9), we can connect C to the mutual information $I(x, x^+)$:

$$C = \frac{\mathbb{E}_{(x, x^+) \sim p(x, x^+)} [f(x, x^+)]}{I(x, x^+)}. \quad (13)$$

Proof of Lemma 4. We can first consider \mathcal{J}_1 without loss of generality. Denoting $g(x, x^+) = e^{f(x, x^+)}$, we take the functional derivative:

$$\begin{aligned} \delta \mathcal{J}_1(\log g) &= \lim_{B \rightarrow \infty} \delta \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\frac{1}{B} \sum_{i=1}^B \log \frac{g(x_i, x_i^+)}{\sum_{j \neq i} g(x_i, x_j^+)} \right] \\ &= \lim_{B \rightarrow \infty} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\frac{1}{B} \sum_{i=1}^B \frac{(\sum_{j \neq i} g(x_i, x_j^+)) \delta g(x_i, x_i^+) - g(x_i, x_i^+) \delta (\sum_{j \neq i} g(x_i, x_j^+))}{g(x_i, x_i^+) (\sum_{j \neq i} g(x_i, x_j^+))} \right] \end{aligned}$$




Figure 7: Our approach enables a goal-conditioned policy to reach farther targets (red) from the start (green) by planning over intermediate waypoints (orange).

$$\begin{aligned}
&= \lim_{B \rightarrow \infty} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\frac{1}{B} \sum_{i=1}^B \frac{\delta g(x_i, x_i^+)}{g(x_i, x_i^+)} - \frac{\delta(\sum_{j \neq i} g(x_i, x_j^+))}{\sum_{j \neq i} g(x_i, x_j^+)} \right] \\
&= \lim_{B \rightarrow \infty} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\frac{1}{B} \sum_{i=1}^B \int \left(\left(\frac{\delta g(x_i, x_i^+)}{g(x_i, x_i^+)} \right) p(x^+ | x_i) \right. \right. \\
&\quad \left. \left. - \sum_{k \neq i} \left(\frac{\delta g(x_i, x_i^+)}{g(x_i, x_i^+) - g(x_i, x_k^+) + \sum_{j \neq i} g(x_i, x_j^+)} \right) p(x^+) \right) dx^+ \right] \\
&= \lim_{B \rightarrow \infty} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\frac{1}{B} \sum_{i=1}^B \int \delta g(x_i, x_i^+) \left(\frac{p(x^+ | x_i)}{g(x_i, x_i^+)} \right. \right. \\
&\quad \left. \left. - \underbrace{\mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B} \left[\frac{1}{\sum_{j \neq i} g(x_i, x_j^+)} \right]}_{\text{as } B \rightarrow \infty} \right) p(x^+) dx^+ \right] \\
&= \lim_{B \rightarrow \infty} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\frac{1}{B} \sum_{i=1}^B \int \delta g(x_i, x_i^+) \left(\frac{p(x^+ | x_i)}{g(x_i, x_i^+)} \right. \right. \\
&\quad \left. \left. - \underbrace{\mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B} \left[\frac{1}{\sum_{j \neq i} g(x_i, x_j^+)} \right]}_{\triangleq k(x_i) \text{ indep. of } x^+} \right) p(x^+) dx^+ \right] \\
&= \lim_{B \rightarrow \infty} \mathbb{E}_{\{(x_i, x_i^+)\}_{i=1}^B \sim p(x, x^+)} \left[\frac{1}{B} \sum_{i=1}^B \int \delta g(x_i, x_i^+) \left(\frac{p(x^+ | x_i)}{g(x_i, x_i^+)} - k(x_i) p(x^+) \right) dx^+ \right] \\
&= \int \delta g(x, x^+) \left(\frac{p(x^+ | x)}{g(x, x^+)} - k(x) p(x^+) \right) dx^+.
\end{aligned}$$

This is zero when

$$g(x, x^+) = \frac{p(x | x^+)}{k(x)p(x)},$$

i.e.,

$$f(x, x^+) = \log p(x | x^+) - \log \underbrace{c_1(x)}_{k(x)p(x)}.$$

as in Eq. (10), and Eq. (11) follows similarly, exchanging x and x^+ . \square

B Additional Experiments

Fig. 7 visualizes the inferred waypoints from the task in Fig. 5.

Fig. 8 visualizes the representations learned on a 46-dimensional robotic hammering task (see Section 5.3).




Figure 8: Planning for 46-dimensional robotic hammering. *(Left)* A dataset of trajectories demonstrating a hammer knocking a nail into a board [31]. *(Center)* We visualize the learned representations as blue circles, with the transparency indicating the index of that observation along the trajectory. We also visualize the inferred plan (Section 4.3) as red circles connected by arrows. *(Right)* Representations learned by PCA on the same trajectory as *(a, left)*.

B.1 Stock Prediction

We show results on a stock opening price task in Fig. 9.

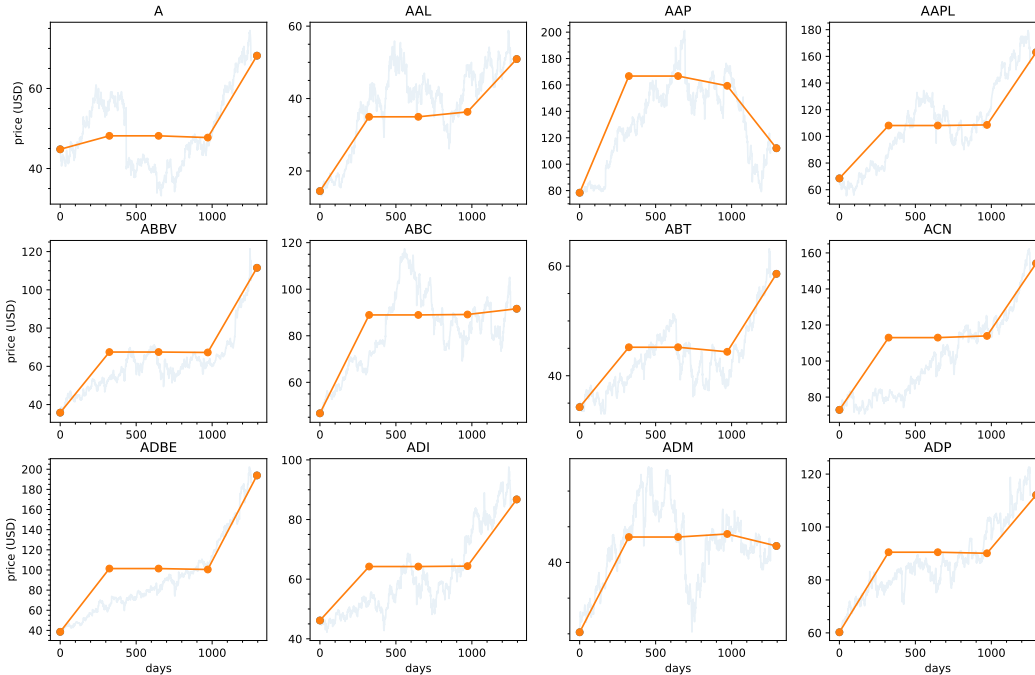


Figure 9: **Stock Prediction.** We apply temporal contrastive learning to time series data of the stock market. Data are the opening prices for the 500 stocks in the S&P 500, over a four year window. We remove 30 stocks that are missing data. For evaluation, we choose a 100 day window from a validation set, and use Theorem 4.1 to perform “inpainting”, predicting the intermediate stock prices *jointly* for all stocks (orange), given the first and last stock price. The true stock prices are shown in blue. While we do not claim that this is a state-of-the-art model for stock prediction, this experiment demonstrates another potential application of our theoretical results.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The claims in the abstract and introduction are supported by the Lemmas in the main text.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: See "Limitations" in Sec. 6, and discussion following the Assumptions in the main text.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Proofs for each Lemma are provided in either the main text or the appendices.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All results and figures in paper can be reproduced by running make in the linked code.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Code has been released. The datasets used have been released by prior work.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Details are in the appendix and/or the supplemental code.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: See Sec. 5.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: See Sec. 5

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: We are unaware of any violations.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This paper provides theoretical analysis of representations learned by contrastive learning. This analysis suggests new ways of using these representations, for both beneficial and harmful purposes.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: N/A

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: See citations to the D4RL dataset.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [No]

Justification: N/A

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: N/A

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: N/A

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.